

NVM Express Technical Errata

Errata ID	003
Change Date	4/4/2013
Affected Spec Ver.	NVM Express 1.0 and 1.1
Corrected Spec Ver.	

Submission info

Name	Company	Date
Amber Huffman	Intel	2/7/2013
Eric Peterson	Synopsys	2/7/2013
Judy Brock	Samsung	2/7/2013
Frank Chu	HGST	2/8/2013
John Carroll	Intel	2/8/2013
Dave Landsman	SanDisk	2/8/2013

Updated the error log to describe the correct settings when the error is not specific to a particular command.

Clarified that the time limited error recovery feature applies only to error recovery.

Updated the Metadata SGL pointer to avoid issues with unknown size of entry.

Updated volatile write cache behavior.

Description of the specification technical flaw:

Modify section 5.10.1.1 as shown below:

This log page is used to describe extended error information for a command that completed with error **or report an error that is not specific to a particular command**. Extended error information is provided when the More (M) bit is set to '1' in the Status Field for the completion queue entry associated with the command that completed with error or as part of an asynchronous event with an Error status type. This log page is global to the device.

This error log may return the last n errors. If host software specifies a data transfer of the size of n error logs, then the error logs for the last n errors is returned. The ordering of the entries is based on the time when the error occurred, with the most recent error being returned as the first log.

Each entry in the log page returned is defined in Figure 74. The log page is a set of 64 byte entries; the number of entries supported is indicated in the Identify Controller data structure in Figure 82.

Figure 74: Get Log Page – Error Information Log Entry (Log Identifier 01h)

Bytes	Description								
07:00	Error Count: This is a 64-bit incrementing error count, indicating a unique identifier for this error. The error count starts at 1h, is incremented for each unique error log entry, and is retained across power off conditions. A value of 0h indicates an invalid entry; this value may be used when there are lost entries or when there are fewer errors than the maximum number of entries the controller supports.								
09:08	Submission Queue ID: This field indicates the Submission Queue Identifier of the command that the error information is associated with. If the error is not specific to a particular command then this field shall be set to FFFFh.								
11:10	Command ID: This field indicates the Command Identifier of the command that the error is associated with. If the error is not specific to a particular command then this field shall be set to FFFFh.								
13:12	Status Field: This field indicates the Status Field for the command that completed. The Status Field is located in bits 15:01, bit 00 corresponds to the Phase Tag posted for the command. If the error is not specific to a particular command then this field reports the most applicable status value.								
15:14	<p>Parameter Error Location: This field indicates the byte and bit of the command parameter that the error is associated with, if applicable. If the parameter spans multiple bytes or bits, then the location indicates the first byte and bit of the parameter.</p> <table border="1"> <thead> <tr> <th>Bits</th><th>Description</th></tr> </thead> <tbody> <tr> <td>7:0</td><td>Byte in command that contained the error. Valid values are 0 to 63.</td></tr> <tr> <td>10:8</td><td>Bit in command that contained the error. Valid values are 0 to 7.</td></tr> <tr> <td>15:11</td><td>Reserved</td></tr> </tbody> </table> <p>If the error is not specific to a particular command then this field shall be set to FFFFh.</p>	Bits	Description	7:0	Byte in command that contained the error. Valid values are 0 to 63.	10:8	Bit in command that contained the error. Valid values are 0 to 7.	15:11	Reserved
Bits	Description								
7:0	Byte in command that contained the error. Valid values are 0 to 63.								
10:8	Bit in command that contained the error. Valid values are 0 to 7.								
15:11	Reserved								
23:16	LBA: This field indicates the first LBA that experienced the error condition, if applicable.								
27:24	Namespace: This field indicates the namespace that the error is associated with, if applicable.								
28	Vendor Specific Information Available: If there is additional vendor specific error information available, this field provides the log page identifier associated with that page. A value of 00h indicates that no additional information is available. Valid values are in the range of 80h to FFh.								
63:29	Reserved								

Modify Figure 96 as shown below:

Figure 96: Error Recovery – Command Dword 11

Bit	Description
31:16	Reserved
15:00	Time Limited Error Recovery (TLER): Indicates a limited retry timeout value in 100 millisecond units. This applies to I/O (e.g. Read, Write, etc) commands that indicate a time limit is required. The timeout starts when error recovery actions have started while processing the command. A value of 0h indicates that there is no timeout. Note: This mechanism is primarily intended for use by host software that may have alternate means of recovering the data.

Modify the fourth paragraph of section 6 as shown below:

In the case of Compare, Read, **and Write, and Write Zeroes** commands, the host may indicate whether a time limit should be applied to **error recovery for** the operation by setting the Limited Retry (LR) field in the command. The time limit is specified in the Error Recovery feature, specified in section 5.12.1.5. If the host does not specify a time limit should be applied, then the controller should apply all error recovery means to complete the operation regardless of the total time required.

Modify Figure 12 in section 4.2 as shown below:

Figure 12: Command Format – NVM Command Set

Bytes	Description		
23:16	<p>If CDW0[15] is cleared to '0', then the definition of this field is:</p> <table><tr><td>23:16</td><td>Metadata Pointer (MPTR): This field contains the address of a contiguous physical buffer of metadata. This field is only used if metadata is not interleaved with the logical block data, as specified in the Format NVM command. This field shall be Dword aligned.</td></tr></table>	23:16	Metadata Pointer (MPTR): This field contains the address of a contiguous physical buffer of metadata. This field is only used if metadata is not interleaved with the logical block data, as specified in the Format NVM command. This field shall be Dword aligned.
	23:16	Metadata Pointer (MPTR): This field contains the address of a contiguous physical buffer of metadata. This field is only used if metadata is not interleaved with the logical block data, as specified in the Format NVM command. This field shall be Dword aligned.	
<p>If CDW0[15] is set to '1', then the definition of this field is:</p> <table><tr><td>23:16</td><td>Metadata SGL Segment Pointer (MSGLP): This field contains the address of an SGL segment containing exactly one SGL Descriptor which describes the metadata to transfer. This field is only used if metadata is not interleaved with the logical block data, as specified in the Format NVM command. This field shall be Qword aligned. Refer to section 4.4.</td></tr></table>	23:16	Metadata SGL Segment Pointer (MSGLP): This field contains the address of an SGL segment containing exactly one SGL Descriptor which describes the metadata to transfer. This field is only used if metadata is not interleaved with the logical block data, as specified in the Format NVM command. This field shall be Qword aligned. Refer to section 4.4.	
23:16	Metadata SGL Segment Pointer (MSGLP): This field contains the address of an SGL segment containing exactly one SGL Descriptor which describes the metadata to transfer. This field is only used if metadata is not interleaved with the logical block data, as specified in the Format NVM command. This field shall be Qword aligned. Refer to section 4.4.		

Modify a portion of section 6.3 as shown below:

After a write has completed, reads for that location which are subsequently submitted shall return the data from that write and not an older version of the data from a previous write **with the following exception;**

If all of the following conditions are met:

- a) the controller supports a volatile write cache;
- b) the volatile write cache is enabled;
- c) the FUA bit for the write is not set;
- d) no flush commands, associated with the same namespace as the write, successfully completed before shutdown; and

e) a controller shutdown occurs without completing the normal or abrupt shutdown procedure outlined in Section 7.6.2

then subsequent reads for locations written to the volatile write cache that were not written to non-volatile storage may return older data.

Disposition log

2/7/2013	Erratum captured.
2/8/2013	Updates based on 2/7 meeting.
4/3/2013	Added volatile write cache clarification agreed to in 3/21 meeting
4/4/2013	Editorial change to make exception conditions stand out. Approved at 4/4 Workgroup meeting for 30 day review
5/16/2013	Erratum ratified.

Technical input submitted to the NVMHCI Workgroup is subject to the terms of the NVMHCI Contributor's agreement.